

CS 433: Computer Architecture – Fall 2021

Homework 7

Total Points: Undergraduates (29 points), Graduates (42 points)

Undergraduate students should only solve the first 4 problems.

Graduate students should solve all problems.

Due Date: December 8, 2021 at 10:00 pm CT

(See course information slides for more details)

Directions:

- All students must write and sign the following statement at the end of their homework submission. "I have read the honor code for this class in the course information handout and have done this homework in conformance with that code. I understand fully the penalty for violating the honor code policies for this class." No credit will be given for a submission that does not contain this signed statement.
- On top of the first page of your homework solution, please write your name and NETID, your partner's name and NETID, and whether you are an undergrad or grad student.
- Name your homework solution file as *firstname_lastname_hw7.pdf*
- Please show all work that you used to arrive at your answer. Answers without justification will not receive credit. Errors in numerical calculations will not be penalized. Cascading errors will usually be penalized only once.
- See course information slides for more details.

Problem 1 [13 points]

This problem concerns MOESI, an invalidation based snooping cache coherence protocol, for bus-based shared-memory multiprocessors with a single level of cache per processor. The MOESI protocol has five states. A block starting at address *Addr* can be in one of the following states in cache C:

- *Modified*: The block is present only in cache C and the data in the cache is dirty or modified (i.e., it reflects a more recent version than the copy in memory).
- *Owned*: The block is present in cache C and may also be present in other caches. The memory may not have an up-to-date copy of this block. The block is said to be owned by cache C and C must service the requests of other caches to this block since memory may not have an up-to-date copy.
- *Exclusive*: The block is present in a single cache (C) but is clean (i.e., memory has an up-to-date copy of the block).
- *Shared*: The block is present in cache C and possibly present in other caches.
- *Invalid*: The block is not valid in cache C (space for the block may or may not be currently allocated in this cache).

If the cache has a block in Owned state, then it services any requests to that block from other processors. Assume that the memory does NOT update its copy even if the request is a read by some other cache and the Owner cache has put the block on the bus. Hence, the owner cache remains in Owned state and continues to service other requests, until the block is replaced from its cache. Also assume that the only way to reach the Owned state is from the Modified or Exclusive state, when some other cache issues a read request for that block.

If a cache C has a block in exclusive or modified state, then it is responsible for servicing any requests to that block from other processors. If the request is a read, then the cache C transitions to the Owned state, and memory does NOT update its copy.

On replacement of a block in Owned or Modified state, the block is sent to memory, and memory resumes responsibility for servicing subsequent requests to that block. Replacement of a block in Exclusive state is similar, except that the block need not be sent to memory (since memory already has a copy).

Assume that after a cache performs a transaction on the bus, there is a mechanism for it to know whether other caches have a copy of the requested block or not at that time. This enables the cache to determine whether to transition to exclusive state.

Part A [2 points]

Consider a Block B in Owned state in the cache of processor P. Can B be in a non-invalid state in any other processor's cache? If yes, then what are the possible (non-invalid) states in which B could be in any of the other caches? If no, then explain why not.

Solution:

Yes; B can be in Shared state in other caches even when it is in Owned state in P. This scenario results when P has the block in Modified state and another processor does a read on B. The only possible non-invalid state of B in other caches is S.

Grading:

1 point for recognizing that B may be non-invalid in another processor's cache.

1 point for the correct possible state.

Part B [6 points]

This part concerns the response of the cache of processor i to bus transactions initiated by the cache of processor j for a block that starts at address B (referred to as block B below). Fill out the following rows for the state transition table for the cache of processor i , showing the next state for block B in the cache and any action taken by the cache. Each entry should be filled out as: *Next State/Action* (e.g., S/Send block to memory) where

Next State = **M, O, E, S, or I**

Action = **Send block to memory, Send block to cache, Send block to cache and memory, or No action**

Note: If an entry is not possible (i.e., the system cannot be in such a state), write “Not Possible”.

Current state of block B in cache of processor i	Read of block B by cache of processor j	Invalidate of block B by cache of processor j (with no read request for the block)	Read + Invalidate of block B by cache of processor j
M	O/Send block to j	Not Possible	I/Send Block to j
O	O/Send block to j	I/No Action	I/Send Block to j

Grading:

1 point for each entry.

Partial credit is awarded if at least the next state is computed correctly.

½ point if the answer is “I/send block to j and memory” in both rightmost cells.

Part C [5 points]

Consider the following sequence of operations by two processors for a block that starts at address B. Determine the state of that block in the caches of both the processors after each operation in the sequence for the MOESI protocol. Both caches are initially empty and all lines are in the I state. The table below is provided to help organize your answer.

No	Operation	MOESI	
		P1	P2
1	P1 reads B	E*	I
2	P1 writes B	M	I
3	P2 writes B	I	M
4	P1 reads B	S	O
5	P1 writes B	M	I
6	P2 reads B	O	S

*Since P1 performed a read and has the only cached copy, it will transition to E instead of S.

Grading:

½ point per entry

Problem 2 [6 points]

compare-and-swap($R1, R2, L$) is an atomic synchronization primitive which atomically compares the value in memory location L with $R1$, and if and only if they are equal, exchanges the values in $R2$ and L . *compare-and-swap*($R1, R2, L$) can be used to efficiently emulate many other primitives.

Part A [2 points]

Implement an atomic *test-and-set* on memory address L in assembly using *compare-and-swap*($R1, R2, L$) as the only atomic primitive. Let $L = 1$ when the lock is taken and $L = 0$ when it is free, and these will be the only values present at L . You may use any registers you like.

Solution:

```
DADDI R1, R0, R0
DADDI R2, R0, #1
CAS R1, R2, 0(L)
```

(Note that the R0 register is a constant 0.)

Grading:

2 points for a correct solution.

Part B [2 points]

Implement the *test-and-test-and-set* semantics on memory address L in assembly using *compare-and-swap* as the only atomic primitive. Let $L = 1$ when the lock is taken, and $L = 0$ when it is free. You can use any registers you like as well as ordinary loads and stores. Include any instructions needed to ensure that the operation eventually completes successfully, as if you are actually trying to acquire a lock.

Solution:

```
LOCK: DADDI R2, R0, #1
LD R1, 0(L)
BNEZ R1, LOCK
CAS R1, R2, 0(L)
BNEZ R2, LOCK
```

(Note that the R0 register is a constant 0.)

Grading:

2 points for a correct solution.

1 point if the CAS is used correctly, but the loop back branch on failure is not included or is incorrect.

Part C [2 points]

Use compare-and-swap to implement an atomic *fetch-and-increment*(*R1*, *L*) in assembly, which atomically copies the old value in *L* to *R1* and then increments the value in *L* by 1. Again, you can use any registers you like well as ordinary loads and stores. Include any instructions needed to ensure that the operation eventually completes successfully; i.e., the increment must be guaranteed to occur atomically.

Solution:

```
FAA: LD      R1, 0(L)
      DADDI  R2, R1, #1
      CAS   R1, R2, 0(L)
      BNE  R1, R2, FAA
```

Grading:

2 points for a correct solution.

1 point if the CAS is used correctly but the loop back branch on failure is not included or is incorrect.

Problem 3 [4 points]

For this problem, assume sequential consistency. Consider the following code fragments executed on two processors:

Initially $P = Q = R = S = 0$

P1	P2
$P = 5$	$L = Q$
$Q = 12$	$M = P$
$R = 8$	$N = S$
$S = 6$	$O = R$

Part A [2 points]

The processors execute instructions independent of each other. Thus, the order of execution of instructions cannot be determined *a priori* if no constraint is placed on the execution of the instructions. What synchronization is required to ensure that all of P1's instructions are to be executed before any of P2's instructions, as given above? Ensure your solution makes clear the constituent memory operations used for the synchronization; i.e., don't insert just a call to a function or library without the code for that function/library. Unnecessarily inefficient solutions will not get full credit.

Solution:

Initially $P = Q = R = S = 0$

Initially $V = 0$

P1	P2
	<code>while (V == 0);</code>
$P = 5$	$L = Q$
$Q = 12$	$M = P$
$R = 8$	$N = S$
$S = 6$	$O = R$
$V = 1$	

Grading:

2 points for correct solution.

1 point partial credit for barrier after last instruction of P1 and before first instruction of P2.

Part B [2 points]

Suppose we don't care about the relative order of execution of the sets of instructions, but we want atomicity in execution; i.e., we want all instructions of one processor to complete before any instructions are executed in the other processor. What synchronization is required to ensure this?

Solution:

This can be achieved by placing the code fragments between lock and unlock operations. These operations can be implemented using Test & Set or any other technique but must all be to the same location.

P1	P2
Lock(x)	Lock(x)
P = 5	L = Q
Q = 12	M = P
R = 8	N = S
S = 6	O = R
Unlock(x)	Unlock(x)

Grading:

2 points for correct solution.

Problem 4 [6 points]

This problem depends on your solution to part A of problem 3. You will get credit for this problem only if part A of problem 3 is correct.

Part A [3 points]

On a sequentially consistent system, what values can L, M, N, and O have at the end of execution for the code in your solution of problem 3, part A? You must explain your answer for full credit.

Solution:

$L = 12$, $M = 5$, $N = 6$, $O = 8$. On a sequentially consistent processor, we can assume that P2's reads of P, Q, R, and S won't occur until its while loop terminates; i.e., until P2 sees the value 1 for V. Once P2 sees the value 1 for V, we know that P1's write to V has occurred and so we can assume that P1's writes to P, Q, R, and S are complete. We are therefore guaranteed that P2's reads of P, Q, R, and S will see the values written to those variables by P1.

Grading:

1 point for correct values of L, M, N, and O.

2 points for the correct explanation.

Part B [3 points]

On a system that is not known to be sequentially consistent, what possible values can variables L, M, N, and O have at the end of execution, for the code in your solution of problem 3 part A? You must explain your answer for full credit.

Solution:

$L = 0$ or 12 , $M = 0$ or 5 , $N = 0$ or 6 , $O = 0$ or 8 . Since the system is not known to be sequentially consistent, it is possible that some or all of P2's reads of P, Q, R, S actually occur before its while loop terminates, or some or all of P1's writes to P, Q, R, S actually occur after P1's write to V. So P2 is not guaranteed to see P1's writes to P, Q, R, S, and may see only a subset of those writes.

Grading:

2 points for listing all possible correct values of L, M, N, and O; 1 point if only the values 0 are listed.

1 point for a correct explanation which should include some mention of how the value 0 could be obtained for any of these variables.

NOTE: ONLY GRADUATE STUDENTS SHOULD SOLVE THIS PROBLEM

Problem 5 [13 points]

In the discussion of the barrier semantics in the lecture, a common parallel programming pattern was described where all processors produce data during one phase of a computation, and that data is consumed by many (or all) processors during the next phase. The “barrier” is used as synchronization so that the consumers of the data in phase $n+1$ are guaranteed to see the data that was produced in phase n .

An example program with such a pattern is molecular dynamics simulation. The outermost loop iterates over discrete timesteps covering the period of time covered by the simulation. In each timestep, we need to do two things:

(1) For each particle (or molecule) in the simulated system, we compute the force on the particle exerted by every other particle in the system. The force between two particles is a function of the distance between the two particles (i.e., a function of the positions of the two particles, as computed during the previous time step - see below). The total force on a given particle is the sum of the forces generated by all particles within the cutoff distance of the given particle.

(2) After the force on a particle is computed, its new position is computed using the force computed above and its old position computed in the previous timestep. The new positions are then used as input for the next timestep.

Part A [5 points]

How would you parallelize the computation for each timestep? Specifically, indicate how the computation is divided among N processors, which shared-memory variables are read and written by a given processor in different phases of the program, and where does what synchronization appear in the program and for what purpose. You can describe your algorithm in words and do not have to give code.

Solution: Assume that there are P particles, with $P \geq N$. Assign P/N particles for which to compute forces and positions to each processor. Each timestep consists of two phases on each processor:

1. For each particle p that it has been assigned, the processor reads the positions of all other particles and computes the force exerted on p .
2. For each particle p that it has been assigned, the processor reads the computed force on p and computes its new position.

A barrier is required at the end of each phase. The first barrier ensures that all processors read all the old positions before the positions are updated in the next phase. The second barrier ensures that all the positions are updated before they are read in the next timestep to calculate the forces on other particles.

Grading:

1 point for dividing particles among the processors, assigning P/N particles per processor.

1 point for the first phase of each timestep, computing forces.

1 point for the second phase of each timestep, computing new positions.

2 points for the barriers after each phase and an adequate explanation of why they are needed.

Part B [8 points]

Consider a shared-memory multiprocessor system where each processor has a private L1 cache and all processors and memory are connected by a bus. Suppose the L1 caches are magically of infinite size. Assume an MSI cache coherence protocol. Consider the 20th timestep in the simulation. What will be the state of all the shared variables in a processor's cache at the end of each synchronization event in this timestep? Which accesses will result in bus traffic in this timestep?

Solutions:

1. At the end of the first barrier:
 - The forces of all particles assigned to a processor will be in the M state.
 - All positions of all particles will be in the S state.
2. At the end of the second barrier:
 - The forces of all particles assigned to a processor will remain in the M state.
 - The positions of all particles assigned to a processor will be in the M state.
 - The positions of all *other* particles will be in the I state.

In the first phase, the reads of the positions of all particles assigned to other processors will result in misses. In the second phase, the writes to the positions of the particles will result in invalidations.

Grading:

2 points for correct state after the first barrier.

2 points for correct state after the second barrier.

2 points for correct bus traffic in the first phase.

2 points for correct bus traffic in the second phase.

If a different correct algorithm is used in part A, then give appropriate credit, allocating:

4 points for correct state after synchronizations.

4 points for correct bus traffic.

If the algorithm in part A is incorrect, then assign partial credit on a case by case basis, e.g., 0 credit if the algorithm in Part A is completely broken.